

Nov. 9. 2007 4:03PM

pending 28-36, 38-41

09/927,589  
No. 4004 P. 9

cancelled: 1-27, 37,

**LISTING OF CLAIMS:**

This listing of claims will replace all prior versions, and listings, of claims in the application:

Claims 1-27. (Canceled)

28. (Currently Amended) A method of minimizing network traffic during the process of synchronizing a local transaction database with a remote transaction database, comprising:

receiving from the remote transaction database a single summary hash computed for database records lying in an interval in the remote transaction database;

comparing the single summary hash to a local single summary hash computed for database records lying in a same interval in the local transaction database; and

when the single summary hash does not match the local summary hash, partitioning the interval into at least a first and a second sub-interval and requesting from the remote transaction database a first sub-interval summary hash and a second sub-interval summary hash, the first sub-interval summary hash computed for database records lying in the first sub-interval in the remote transaction database and the second sub-interval summary hash computed for database records lying in the second sub-interval in the remote transaction database,

whereby the database can avoid ~~synchronization of~~ updating database records in a sub-interval at the occurrence of a match between with a local sub-interval summary hash ~~that matches and~~ a sub-interval summary hash received from the remote transaction database.

29. (Previously Presented) The method of claim 28 further comprising the step of: when the local sub-interval summary hash does not match the sub-interval summary hash received from the remote transaction database, further partitioning the sub-interval into at least a first and a second sub-sub-interval and requesting from the remote transaction database a first and second sub-sub-interval summary hash, the first sub-sub-interval summary hash computed for database records lying in the first sub-sub-interval in the remote transaction database and the second sub-sub-interval summary hash computed for database records lying in the second sub-sub-interval in the remote transaction database.

30. (Previously Presented) The method of claim 28 further comprising the step of: when the local sub-interval summary hash does not match the sub-interval summary hash received from the remote transaction database, requesting from the remote transaction database a hash value for each database record lying in the sub-interval.

31. (Previously Presented) The method of claim 28, wherein each database record has a hash value and wherein each summary hash for an interval is computed by combining the hash values for each database record in the interval.

32. (Previously Presented) The method of claim 31, wherein the summary hashes for different intervals are stored in a tree structure.

33. (Previously Presented) The method of claim 32, wherein the tree structure is a B+ tree structure.

34. (Previously Presented) The method of claim 32, wherein the hash values are message digests and wherein the hash values are combined to compute a summary hash by an exclusive or (XOR) of the hash values.

35. (Currently Amended) A computer implemented synchronizable transactional database comprising:

a database;

an interval hash value computing module coupled to the database and configured to compute a summary hash of a plurality of hash values, each hash value associated with a database record lying in an interval of the database; and

a synchronization module coupled to the database and to the interval hash value computing module, the synchronization module configured to identify database records that need synchronization by comparing a summary hash from the interval hash value computing module computed for database records lying in an interval of the database with a remote summary hash received from a remote transaction database, the synchronization module further configured to partition an interval into at least a first and a second sub-interval when a summary hash for the interval in the database does not match a remote summary hash so as to seek remote summary hashes for the first and second sub-intervals from the remote transaction database,

whereby the database can avoid ~~synchronization of updating~~ database records in a sub-interval at the occurrence of a match between with a local sub-interval summary hash ~~that matches and~~ a sub-interval summary hash received from the remote transaction database.

36. (Previously Presented) The system of claim 35, wherein the database further comprises a transactional support layer configured to support a storage layer of the database, the transactional support layer further comprising shadow blocks which provides for atomized updates to the storage layer.

37. (Canceled)

38. (Previously Presented) The system of claim 36, wherein the interval hash value computing module computes a summary hash for database records lying in an interval in the database by combining the hash values associated with each database record lying in the interval.

39. (Previously Presented) The system of claim 38, wherein the summary hashes for different intervals are stored in a tree structure.

40. (Previously Presented) The system of claim 39, wherein the tree structure is a B+ tree structure.

41. (Previously Presented) The system of claim 39, wherein the hash values are message digests and wherein the hash values are combined to compute a summary hash by using an exclusive or (XOR) of the hash values.